arXiv:1701.08785v1 [cs.IT] 21 Jan 2017

Universal Tree Source Coding Using Grammar-Based Compression*

Danny Hucke and Markus Lohrey

Universität Siegen, Germany

February 1, 2017

Abstract

We apply so-called tree straight-line programs to the problem of universal source coding for binary trees. We derive an upper bound on the maximal pointwise redundancy (or worst-case redundancy) that improve previous bounds on the average case redundancy obtained by Zhang, Yang, and Kieffer using directed acyclic graphs. Using this, we obtain universal codes for new classes of tree sources.

1 Introduction

Universal source coding for finite sequences over a finite alphabet Σ (i.e., strings over Σ) is a well-established topic of information theory. Its goal is to find prefix-free lossless codes that are universal (or optimal) for classes of information sources. In a series of papers, Cosman, Kieffer, Nelson, and Yang developed grammar-based codes that are universal for the class of finite state sources [8, 9, 10, 16]. Grammar-based compression works in two steps: In a first step, from a given input string $w \in \Sigma^*$ a context-free grammar $\mathcal{G}_w$ that produces only the string w is computed. Context-free grammars that produce exactly one string are also known as *straight-line programs*, briefly SLPs, and are currently an active topic in text compression and algorithmics on compressed texts, see [12] for a survey. In a second step, the SLP $\mathcal{G}_w$ is encoded by a binary string $B(\mathcal{G}_w)$. There exist several algorithms that compute from a given input string w of length n an SLP $\mathcal{G}_w$ of size $O(n/\log n)$ (the size of an SLP is the total number of symbols in all right-hand sides of the grammar) [8]; the best known example is probably the LZ78 algorithm [18]. By combining any of these algorithms with the

*This work has been supported by the DFG research project LO 748/10-1 (QUANTKOMP).

binary encoder B for SLPs from [8], one obtains a grammar-based encoder $E : \Sigma^* \to \{0,1\}^*$, whose worst case redundancy for input strings of length n is bounded by $O(\log\log n/\log n)$ for every finite state information source over the alphabet Σ. Here, the worst case redundancy for strings of length n is defined as

$$\max_{w\in\Sigma^n, P(w)>0} n^{-1}\cdot(|E(w)| + \log_2 P(w)),$$

where $P(w)$ is the probability that the finite state information source emits w. Thus, the worst case redundancy measures the maximal additive deviation of the code length from the self information, normalized by the length of the source string.

Over the last few years, we have seen increasing efforts aiming to extend universal source coding to structured data like trees [11, 14, 17] and graphs [2, 7]. In this paper, we are concerned with universal source coding for trees. In their recent paper [17], Kieffer, Yang, and Zhang started to extend their work on grammar-based source coding from strings to binary trees. For this, they first represent the input tree t by its *minimal directed acyclic graph* $\mathcal{D}_t$ (the minimal DAG of t). This is the directed acyclic graph obtained by removing multiple occurrences of the same subtree from t. In a second step, the minimal DAG $\mathcal{D}_t$ is encoded by a binary string $B(\mathcal{D}_t)$; this step is similar to the binary coding of SLPs from [8]. Combining both steps yields a tree encoder $E_{\mathrm{dag}} : \mathcal{T} \to \{0,1\}^*$, where $\mathcal{T}$ denotes the set of all binary trees. In order to define universality of such a tree encoder, the classical notion of an information source on finite sequences is replaced in [17] by the notion of a structured tree source. The precise definition can be found in Section 2.2; for the moment the reader can think about a collection of probability distributions $(P_n)_{n\in\mathbb{N}}$, where every P_n is a distribution on a finite non-empty subset F_n, and these sets partition $\mathcal{T}$. The main cases considered in [17] as well as in this paper are: (i) F_n is the set of all binary trees with n leaves (leaf-centric sources) and (ii) F_n is the set of all binary trees of depth n (depth-centric sources). Then, the authors of [17] introduce two properties on binary tree sources: (i) the domination property (see Section 3, where it is called the weak domination property) and (ii) the representation ratio negligibility property. The latter states that $\sum_{t\in F_n} P_n(t)\cdot|\mathcal{D}_t|/|t|$ (the average compression ratio achieved by the minimal DAG) converges to zero for $n\to\infty$, where the size $|t|$ of the binary tree is defined as its number of leaves. The technical main result of [17] states that for every structured tree source $(P_n)_{n\in\mathbb{N}}$ satisfying the domination property and the representation ratio negligibility property the *average case redundancy*

$$\sum_{t\in F_n, P_n(t)>0} |t|^{-1}\cdot(|E_{\mathrm{dag}}(t)| + \log_2 P_n(t))\cdot P_n(t)$$

converges to zero for $n\to\infty$. Finally, two classes of tree sources having the domination property and the representation ratio negligibility property are

presented in [17]. One is a class of leaf centric sources, the other one is a class of depth centric sources. Both sources have the property that every tree with a non-zero probability is balanced in a certain sense, the precise definitions can be found in Section 3.3 and Section 3.4. As a first contribution, we show that for these sources not only the average case redundancy but also the worst case redundancy

$$\max_{t \in F_n, P_n(t)>0} |t|^{-1} \cdot (|E_{\mathrm{dag}}(t)| + \log_2 P_n(t)) \tag{1}$$

converges to zero for $n \to \infty$. More precisely, we show that (1) is bounded by $O(\log\log n/\log n)$ (respectively, $O((\log\log n)^2/\log n)$) for the presented class of leaf-centric tree sources (respectively, depth-centric tree sources). To prove this, we use results from [5, 6] according to which the minimal DAG of a suitably balanced binary tree of size n is bounded by $O(n/\log n)$, respectively $O(n \cdot \log\log n/\log n)$.

Our second main contribution is the application of *tree straight-line programs*, briefly TSLPs, for universal tree coding. A TSLP is a context-free tree grammar that produces exactly one tree, see Section 2.3 for the precise definition and [13] for a survey. TSLPs can be viewed as the proper generalization of SLPs for trees. Whereas DAGs only have the ability to share repeated subtrees of a tree, TSLPs can also share repeated tree patterns with a hole (so-called contexts). In [5], the authors presented a linear time algorithm that computes for a given binary tree t of size n a TSLP $\mathcal{G}_t$ of size $O(n/\log n)$. This shows the main advantage of TSLPs over DAGs: There exist trees of any size n for which the minimal DAG has size n as well. In Section 4.2 we define a binary encoding B of TSLPs similar to the ones for SLPs [8] and DAGs [17]. We then consider the combined tree encoder $E_{\mathrm{tslp}} : \mathcal{T} \to \{0,1\}^*$ with $E_{\mathrm{tslp}}(t) = B(\mathcal{G}_t)$, and prove that its worst case redundancy (that is defined as in (1) with E_{dag} replaced by E_{tslp}) is bounded by $O(\log\log n/\log n)$ for every structured tree source that satisfies the *strong domination property* defined in Section 4.3. The strong domination property is a strengthening of the domination property from [17], and this is what we have to pay extra for our TSLP-based encoding in contrast to the DAG-based encoding from [17]. On the other hand, our approach has two main advantages over [17]:

- The representation ratio negligibility property from [17] is no longer needed.
- We get bounds on the worst case redundancy instead of the average case redundancy.

Both advantages are based on the fact that the grammar-based compressor from [5] computes a TSLP of worst case size $O(n/\log n)$ for a binary tree of size n.

Finally, we present a class of leaf-centric sources (Section 3.3) as well as a class of depth-centric sources (Section 3.4) having the strong domination property. These classes are orthogonal to the classes considered in [17].

2 Preliminaries

In this section, we introduce some basic definitions concerning information theory (Section 2.1), binary trees (Section 2.2) and tree straight-line programs (Section 2.3). The latter are our key formalism for the compression of binary trees.

With $\mathbb{N}$ we denote the natural numbers including 0. We use the standard O-notation and if b is a constant, then we just write $O(\log n)$ for $O(\log_b n)$. For the unit interval $\{r \in \mathbb{R} \mid 0 \leq r \leq 1\}$ we write $[0, 1]$.

2.1 Empirical distributions and empirical entropy

Let $\overline{a} = (a_1, a_2, \ldots, a_n)$ be a tuple of elements that are from some (not necessarily finite) set A. The *empirical distribution* $p_{\overline{a}} : \{a_1, a_2, \ldots, a_n\} \to \mathbb{R}$ of $\overline{a}$ is defined by

$$p_{\overline{a}}(a) = \frac{|\{i \mid 1 \leq i \leq n,\ a_i = a\}|}{n}.$$

We use this definition also for words over some alphabet by identifying a word $w = a_1 a_2 \cdots a_n$ with the tuple $(a_1, a_2, \ldots, a_n)$. The *unnormalized empirical entropy* of $\overline{a}$ is

$$H(\overline{a}) = -\sum_{i=1}^{n} \log p_{\overline{a}}(a_i).$$

A well-known generalization of Shannon's inequality states that for all real numbers $p_1, \ldots, p_k, q_1, \ldots q_k > 0$, if $\sum_{i=1}^{k} p_i = 1 \geq \sum_{i=1}^{k} q_i$ then

$$\sum_{i=1}^{k} -p_i \log_2(p_i) \leq \sum_{i=1}^{k} -p_i \log_2(q_i)$$

see [1] for a proof. As a consequence, for a tuple $\overline{a} = (a_1, a_2, \ldots, a_n)$ with $a_1, \ldots, a_n \in A$ and real numbers $q(a) > 0$ $(a \in A)$ with $\sum_{i=1}^{n} q(a_i) \leq 1$ we have

$$\sum_{i=1}^{n} -\log_2(p_{\overline{a}}(a_i)) \leq \sum_{i=1}^{n} -\log_2(q(a_i)). \tag{2}$$

2.2 Trees, tree sources, and tree compressors

With $\mathcal{T}$ we denote the set of all binary trees. We identify $\mathcal{T}$ with the set of terms that are built from the binary symbol f and the constant a. Formally, $\mathcal{T}$ is the smallest set such that (i) $a \in \mathcal{T}$ and (ii) if $t_1, t_2 \in \mathcal{T}$ then also $f(t_1, t_2) \in \mathcal{T}$. With $|t|$ we denote the number of occurrences of the constant a in t. This is the number of leaves of t. Let $\mathcal{T}_n = \{t \in \mathcal{T} \mid |t| = n\}$ for $n \geq 1$. The depth $d(t)$ of the tree t is recursively defined by $d(a) = 0$ and $d(f(t_1, t_2)) = \max\{d(t_1), d(t_2)\} + 1$. Let $\mathcal{T}^d = \{t \in \mathcal{T} \mid d(t) = d\}$ for $d \in \mathbb{N}$.

Occasionally, we will consider a binary tree t as a graph with nodes and edges in the usual way. Note that a tree $t \in \mathcal{T}_n$ has $2n-1$ nodes in total: n leaves and $n-1$ internal nodes. For a node v we write $t[v]$ for the *subtree* rooted at v in t.

A *context* is a binary tree t, where exactly one leaf is labelled with the special symbol x (called the *parameter*); all other leaves are labelled with a. For a context t we define $|t|$ to be the number of a-labelled leaves of t (which is the number of leaves of t minus 1). We denote with $\mathcal{C}$ the set of all contexts and define $\mathcal{C}_n = \{t \in \mathcal{C} \mid |t| = n\}$ for $n \in \mathbb{N}$. For a tree or context $t \in \mathcal{T} \cup \mathcal{C}$ and a context $s \in \mathcal{C}$, we denote by $s(t)$ the tree or context which results from s by replacing the parameter x by t. For example $s = f(a, x)$ and $t = f(a, a)$ yields $s(t) = f(a, f(a, a))$. The depth $d(t)$ of a context $t \in \mathcal{C}$ is defined as the depth of the tree $t(a)$.

A tree source is a pair $((\mathcal{F}_i)_{i\in\mathbb{N}}, P)$ such that the following conditions hold:

- $\mathcal{F}_i \subseteq \mathcal{T}$ is non-empty and finite for every $i \geq 0$,
- $\mathcal{F}_i \cap \mathcal{F}_j = \emptyset$ for $i \neq j$ and $\bigcup_{i\geq 0} \mathcal{F}_i = \mathcal{T}$, i.e., the sets $\mathcal{F}_i$ form a partition of $\mathcal{T}$,
- $P : \mathcal{T} \to [0, 1]$ and $\sum_{t\in\mathcal{F}_i} P(t) = 1$ for every $i \geq 0$, i.e., P restricted to $\mathcal{F}_i$ is a probability distribution.

In this paper, we consider only two cases for the partition $(\mathcal{F}_i)_{i\in\mathbb{N}}$: either $\mathcal{F}_i = \mathcal{T}_{i+1}$ for all $i \in \mathbb{N}$ (note that there is no tree of size 0) or $\mathcal{F}_i = \mathcal{T}^i$ for all $i \in \mathbb{N}$. Tree source of the former (resp., latter) type are called *leaf-centric* (resp., *depth-centric*).

A *tree encoder* is an injective mapping $E : \mathcal{T} \to \{0, 1\}^*$ such that the range $E(\mathcal{T})$ is prefix-free, i.e., there do not exist $t, t' \in \mathcal{T}$ with $t \neq t'$ such that $E(t)$ is a prefix of $E(t')$. We define the *worst-case redundancy* of E with respect to the tree source $\mathcal{S} = ((\mathcal{F}_i)_{i\in\mathbb{N}}, P)$ as the mapping $i \mapsto R(E, \mathcal{S}, i)$ $(i \in \mathbb{N})$ with

$$R(E, \mathcal{S}, i) = \max_{t\in\mathcal{F}_i, P(t)>0} \frac{1}{|t|} \cdot (|E(t)| + \log_2 P(t))$$

The worst-case redundancy is also known as the *maximal pointwise redundancy*.

2.3 Tree straight-line programs

We now introduce tree straight-line programs for binary trees. Let V be a finite ranked alphabet, where each symbol $A \in V$ has an associated rank 0 or 1. The elements of V are called *nonterminals*. We assume that V contains at least one element of rank 0 and that V is disjoint from the set $\{f, a, x\}$, which are the labels used for binary trees and contexts. We use V_0 (respectively, V_1) for the set of nonterminals of rank 0 (resp. of rank 1). The idea is that nonterminals from V_0 (respectively, V_1) derive to trees from $\mathcal{T}$ (respectively, contexts from $\mathcal{C}$). We denote by $\mathcal{T}_V$ the set of trees over $\{f, a\} \cup V$, i.e. each node in a tree $t \in \mathcal{T}_V$ is labelled with a symbol from $\{f, a\} \cup V$ and the number of children of a node corresponds to the rank of its label. With $\mathcal{C}_V$ we denote the corresponding set of all contexts, i.e., the set of trees over $\{f, a, x\} \cup V$, where the parameter symbol x occurs exactly once and at a leaf position. Formally, we have $\mathcal{T} \subset \mathcal{T}_V$ and $\mathcal{C} \subset \mathcal{C}_V$. A *tree straight-line program* $\mathcal{G}$, or short *TSLP*, is a tuple (V, A_0, r), where $A_0 \in V_0$ is the start nonterminal and $r : V \to (\mathcal{T}_V \cup \mathcal{C}_V)$ is the function which assigns each nonterminal its unique right-hand side. It is required that if $A \in V_0$ (respectively, $A \in V_1$), then $r(A) \in \mathcal{T}_V$ (respectively, $r(A) \in \mathcal{C}_V$). Furthermore, the binary relation $\{(A, B) \in V \times V \mid B \text{ is a label in } r(A)\}$ needs to be acyclic. These conditions ensure that exactly one tree is derived from the start nonterminal A_0 by using the rewrite rules $A \to r(A)$ for $A \in V$. To define this formally, we define $\mathsf{val}_{\mathcal{G}}(t) \in \mathcal{T}$ for $t \in \mathcal{T}_V$ and $\mathsf{val}_{\mathcal{G}}(t) \in \mathcal{C}$ for $t \in \mathcal{C}_V$ inductively by the following rules:

- $\mathsf{val}_{\mathcal{G}}(a) = a$ and $\mathsf{val}_{\mathcal{G}}(x) = x$,
- $\mathsf{val}_{\mathcal{G}}(f(t_1, t_2)) = f(\mathsf{val}_{\mathcal{G}}(t_1), \mathsf{val}_{\mathcal{G}}(t_2))$ for $t_1, t_2 \in \mathcal{T}_V \cup \mathcal{C}_V$ (and $t_1 \in \mathcal{T}_V$ or $t_2 \in \mathcal{T}_V$ since there is at most one parameter in $f(t_1, t_2)$),
- $\mathsf{val}_{\mathcal{G}}(A) = \mathsf{val}_{\mathcal{G}}(r(A))$ for $A \in V_0$,
- $\mathsf{val}_{\mathcal{G}}(A(s)) = \mathsf{val}_{\mathcal{G}}(r(A))(\mathsf{val}_{\mathcal{G}}(s))$ for $A \in V_1$, $s \in \mathcal{T}_V \cup \mathcal{C}_V$ (note that $\mathsf{val}_{\mathcal{G}}(r(A))$ is a context t', so we can built $t'(\mathsf{val}_{\mathcal{G}}(s))$).

The tree defined by $\mathcal{G}$ is $\mathsf{val}(\mathcal{G}) = \mathsf{val}_{\mathcal{G}}(A_0) \in \mathcal{T}$. Moreover, for $A \in V_1$ we also write $\mathsf{val}_{\mathcal{G}}(A)$ for $\mathsf{val}_{\mathcal{G}}(A(x))$.

Example 1. *Let* $\mathcal{G} = (\{A_0, A_1, A_2\}, A_0, r)$ *be a TSLP with* $A_0, A_1 \in V_0, A_2 \in V_1$ *and*

$$r(A_0) = f(A_1, A_2(a)),\ r(A_1) = A_2(A_2(a)),\ r(A_2) = f(x, a).$$

We get $\mathsf{val}_{\mathcal{G}}(A_2) = f(x, a)$, $\mathsf{val}_{\mathcal{G}}(A_1) = f(f(a, a), a)$ *and* $\mathsf{val}(\mathcal{G}) = \mathsf{val}_{\mathcal{G}}(A_0) = f(f(f(a, a), a), f(a, a))$.

In this paper, we will consider two classes of syntactically restricted TSLPs: (i) DAGs (directed acyclic graphs) and (ii) TSLPs in normal form. Let us start with the former; normal form TSLPs will be introduced in Section 4.1.

3 Tree compression with DAGs

In this section we sharpen some of the results from [17], where universal source coding of binary trees using minimal DAGs (directed acyclic graphs) is investigated. In [17], only bounds on the average redundancy for certain classes of tree sources were shown. Here we extend these bounds (for the same classes of tree sources) to the worst-case redundancy.

3.1 Directed acyclic graphs (DAGs)

A DAG is a TSLP $\mathcal{D} = (V, A_0, r)$ such that $V = \{A_0, A_1, \ldots, A_{n-1}\}$ for some $n \in \mathbb{N}$, $n \geq 1$, $V = V_0$ (i.e., all nonterminals have rank 0), and for every $A_i \in V$, the right-hand side $r(A_i)$ is of the form $f(\alpha_1, \alpha_2)$ with $\alpha_1, \alpha_2 \in \{a, A_{i+1}, \ldots, A_{n-1}\}$. Note that a TSLP of this form generates a tree with at least two leaves. In order to include the tree a with a single leaf, we also allow the TSLP $\mathcal{G}_a = (\{A_0\}, A_0, A_0 \mapsto a)$. We define the size of a DAG as $|\mathcal{D}| = n + 1$.

In contrast to general TSLPs, every binary tree t has a unique (up to renaming of nonterminals) minimal DAG $\mathcal{D}_t$, whose size is the number of different (pairwise non-isomorphic) subtrees of t. The idea is to introduce for every subtree $f(t_1, t_2)$ of size at least two a nonterminal A_i with $r(A_i) = f(\alpha_1, \alpha_2)$, where $\alpha_i = a$ if $t_i = a$ and α_i is the nonterminal corresponding to the subtree t_i if $|t_i| \geq 2$ $(i = 1, 2)$. We will only use this minimal DAG $\mathcal{D}_t$ in the sequel.

Example 2. *Consider the tree $t_n = f(f(f(\cdots f(a, a), \cdots a), a), a)$, where f occurs n times. We have $|t_n| = n + 1$. The minimal DAG of t_n is $(\{A_0, A_1, \ldots, A_{n-1}\}, A_0, r_n)$, where $r_n(A_i) = f(A_{i+1}, a)$ for $0 \leq i \leq n - 2$ and $r_n(A_{n-1}) = f(a, a)$ and its size is $n + 1$.*

The above example shows that in the worst-case, the size of the minimal DAG is not smaller than the size of the tree.

3.2 Universal source coding with DAGs

The following condition on a tree source was introduced in [17], where it is called the domination property (later, we will introduce a strong domination property): Let $((\mathcal{F}_i)_{i\in\mathbb{N}}, P)$ be a tree source as defined in Section 2.2. We say that $((\mathcal{F}_i)_{i\in\mathbb{N}}, P)$ has the *weak domination property* if there exists a mapping $\lambda : \mathcal{T} \to \mathbb{R}_{>0}$ with the following properties:

(i) $\lambda(t) \geq P(t)$ for every $t \in \mathcal{T}$.

(ii) $\lambda(f(s, t)) \leq \lambda(s) \cdot \lambda(t)$ for all $s, t \in \mathcal{T}$

(iii) There are constants c_1, c_2 such that $\sum_{t\in\mathcal{T}_n} \lambda(t) \leq c_1 \cdot n^{c_2}$ for all $n \geq 1$.

In [17], the authors define a binary encoding $B(\mathcal{D}_t) \in \{0,1\}^*$, such that $B(\mathcal{D}_t)$ is not a prefix of $B(\mathcal{D}_{t'})$ for all binary trees t, t' with $t \neq t'$. The precise definition of $B(\mathcal{D}_t)$ is not important for us; all we need is the following bound from [17, Theorem 2], where $E_{\mathrm{dag}} : \mathcal{T} \to \{0,1\}^*$ is the tree encoder with $E_{\mathrm{dag}}(t) = B(\mathcal{D}_t)$.

Lemma 1. *Assume that $((\mathcal{F}_i)_{i\in\mathbb{N}}, P)$ has the weak domination property. Let $t \in \mathcal{T}_n$ with $n \geq 2$ and $P(t) > 0$, and let $\mathcal{D}_t$ be the minimal DAG for t. We have*

$$\frac{1}{n} \cdot (|E_{dag}(t)| + \log_2(P(t))) \leq O(|\mathcal{D}_t|/n) + O(|\mathcal{D}_t|/n \cdot \log_2(n/|\mathcal{D}_t|)).$$

This bound is used in [17] to show that for certain leaf-centric and depth-centric tree sources the encoding E_{dag} is universal in the sense that the average redundancy converges to zero. Here, we want to show that for the same tree sources already the worst-case redundancy converges to zero. Let us first define the specific classes of tree sources studied in [17].

3.3 Leaf-centric binary tree sources

We recall the definition of a natural class of leaf-centric tree sources from [17]: Let Σ_{leaf} be the set of all functions $\sigma : (\mathbb{N} \setminus \{0\}) \times (\mathbb{N} \setminus \{0\}) \to [0,1]$ such that for all $n \geq 2$:

$$\sum_{i,j\geq 1,\ i+j=n} \sigma(i,j) = 1. \tag{3}$$

For $\sigma \in \Sigma_{\mathrm{leaf}}$ we define $P_\sigma : \mathcal{T} \to [0,1]$ inductively by:

$$P_\sigma(a) = 1, \tag{4}$$

$$P_\sigma(f(s,t)) = \sigma(|s|,|t|) \cdot P_\sigma(s) \cdot P_\sigma(t). \tag{5}$$

We have $\sum_{t\in\mathcal{T}_n} P_\sigma(t) = 1$ and thus $((\mathcal{T}_i)_{i\geq 1}, P_\sigma)$ is a leaf-centric tree source. The following result is implicitly shown in [17] (see also the proof of Theorem 3).

Lemma 2. *For every $\sigma \in \Sigma_{leaf}$, the leaf-centric tree source $((\mathcal{T}_i)_{i\geq 1}, P_\sigma)$ has the weak domination property.*

We say that a mapping $\sigma \in \Sigma_{\mathrm{leaf}}$ is *leaf-balanced* if there exists a constant c such that for all $(i,j) \in (\mathbb{N} \setminus \{0\}) \times (\mathbb{N} \setminus \{0\})$ with $\sigma(i,j) > 0$ we have

$$\frac{i+j}{\min\{i,j\}} \leq c.$$

In [17] it is shown that for a leaf-balanced $\sigma \in \Sigma_{\mathrm{leaf}}$, the tree source $((\mathcal{T}_i)_{i\geq 1}, P_\sigma)$ has the so called representation ratio negligibility property. This means that the average compression ratio achieved by the minimal DAG (formally, $\sum_{t\in\mathcal{T}_n} P_\sigma(t) \cdot |\mathcal{D}_t|/n$) converges to zero for $n \to \infty$. Using a result from [5], we show the following stronger property.

Lemma 3. *For every leaf-balanced mapping $\sigma \in\in \Sigma_{leaf}$, there exists a constant α such that for every binary tree $t \in \mathcal{T}_n$ with $P_\sigma(t) > 0$ we have $|\mathcal{D}_t| \leq \alpha \cdot n/\log_2 n$.*

Proof. In [5] the authors introduce the notion of β-balanced trees. Let v be a non-leaf node in $t \in \mathcal{T}$, i.e. $t[v] = f(t_1, t_2)$ for binary trees t_1 and t_2. Then v is β-balanced if $|t_1| \leq \beta \cdot |t_2|$ and $|t_2| \leq \beta \cdot |t_1|$. The tree t is β-balanced if for all non-leaf nodes u and v in t such that u is a child node of v, at least one of the nodes u, v is β-balanced. It is shown in [5] that for every constant β there exists a constant α (depending only on β) such that the minimal dag $\mathcal{D}_t$ of a β-balanced tree $t \in \mathcal{T}_n$ has size at most $\alpha \cdot n/\log_2 n$. It follows that we only need to show that a tree $t \in \mathcal{T}_n$ with $P_\sigma(t) > 0$ is β-balanced for a constant β. Since the mapping σ is leaf-balanced, every subtree $f(t_1, t_2)$ of t satisfies $|t_1| + |t_2| \leq c \cdot \min\{|t_1|, |t_2|\}$, where $c \geq 1$ is a constant. Without loss of generality assume that $|t_1| \leq |t_2|$. We get $|t_1| \leq c \cdot |t_2|$ and

$$|t_2| \leq |t_1| + |t_2| \leq c \cdot \min\{|t_1|, |t_2|\} = c \cdot |t_1|,$$

which shows that t is c-balanced. □

Corollary 1. *Let $\sigma \in \Sigma_{leaf}$ be leaf-balanced, and let $\mathcal{S} = ((\mathcal{T}_i)_{i\geq 1}, P_\sigma)$ be the corresponding leaf-centric tree source. Then, we have $R(E_{dag}, \mathcal{S}, i) \leq O(\log\log i/\log i)$.*

Proof. Let α be the constant from Lemma 3. Let $t \in \mathcal{T}_n$ such that $P_\sigma(t) > 0$. Lemma 3 implies $|\mathcal{D}_t| \leq \alpha \cdot n/\log_2 n$. With Lemma 1 and 2 we get

$$\begin{aligned} \frac{1}{n} \cdot (|E_{\mathrm{dag}}(t)| + \log_2(P_\sigma(t))) &\leq O(|\mathcal{D}_t|/n) + O(|\mathcal{D}_t|/n \cdot \log_2(n/|\mathcal{D}_t|)) \\ &\leq O(\log\log n/\log n) \end{aligned}$$

This proves the corollary. □

3.4 Depth-centric binary tree sources

We recall the definition of a natural class of depth-centric tree sources from [17]: Let Σ_{depth} be the set of all mappings $\sigma : \mathbb{N} \times \mathbb{N} \to [0, 1]$ such that for all $n \geq 1$:

$$\sum_{i,j\geq 0,\ \max(i,j)=n-1} \sigma(i, j) = 1. \tag{6}$$

For $\sigma \in \Sigma_{\mathrm{depth}}$, we define $P_\sigma : \mathcal{T} \to [0, 1]$ by

$$P_\sigma(a) = 1, \tag{7}$$

$$P_\sigma(f(s, t)) = \sigma(d(s), d(t)) \cdot P_\sigma(s) \cdot P_\sigma(t). \tag{8}$$

We have $\sum_{t\in\mathcal{T}^n} P_\sigma(t) = 1$ and thus $((\mathcal{T}^i)_{i\geq 0}, P_\sigma)$ is a depth-centric tree source.

The following result is again implicitly shown in [17] and can also be found as a part of the proof of Theorem 4.

Lemma 4. *For every $\sigma \in \Sigma_{depth}$, the depth-centric tree source $((\mathcal{T}_i)_{i\geq 1}, P_\sigma)$ has the weak domination property.*

We say the mapping $\sigma \in \Sigma_{\text{depth}}$ is *depth-balanced* if there exists a constant c such that for all $(i,j) \in \mathbb{N} \times \mathbb{N}$ with $\sigma(i,j) > 0$ we have $|i-j| \leq c$. In [17], the authors define a condition on σ that is slightly stronger than depth-balancedness, and show that for every such σ, the tree source $((\mathcal{T}^i)_{i\geq 0}, P_\sigma)$ has the representation ratio negligibility property. Similarly to Lemma 3, we will show an even stronger property. To do so, we introduce *β-depth-balanced* trees for $\beta \in \mathbb{N}$. A tree t is called *β-depth-balanced* if for each subtree $f(t_1, t_2)$ of t we have $|d(t_1) - d(t_2)| \leq \beta$. Note that for a depth-balanced mapping $\sigma \in \Sigma_{\text{depth}}$, there is a constant β such that every tree t with $P_\sigma(t) > 0$ is β-depth-balanced. We will use the following lemma:

Lemma 5. *Let $\beta \in \mathbb{N}$ and $c = 1 + 1/(1+\beta)$ (thus, $1 < c \leq 2$). For every β-depth-balanced binary tree t, we have $|t| \geq c^{d(t)}$.*

Proof. We prove the lemma by induction on $d(t)$. For the only tree $t = a$ of depth $d(t) = 0$, we have $|t| = 1 = c^0$. Consider now a β-depth-balanced tree $t = f(t_1, t_2)$ of depth $d(t) > 0$. We assume $d(t_1) \geq d(t_2)$, the other case is symmetric. Since t is β-depth-balanced, it follows that $d(t_2) \geq d(t_1) - \beta$. To estimate the size $|t| = |t_1| + |t_2|$, we apply the induction hypothesis to t_1 and t_2, which yields

$$|t| = |t_1| + |t_2| \geq c^{d(t_1)} + c^{d(t_2)} \geq c^{d(t_1)} + c^{d(t_1)-\beta} = c^{d(t_1)} \cdot \left(1 + c^{-\beta}\right).$$

Since $d(t_1) + 1 = d(t)$, it only remains to show that $1 + c^{-\beta} \geq c$, which can be easily done by induction on $\beta \in \mathbb{N}$. □

Lemma 5 together with results from [5, 6] implies:

Lemma 6. *For every depth-balanced mapping $\sigma \in \Sigma_{depth}$ there exists a constant α such that for every binary tree $t \in \mathcal{T}_n$ with $P_\sigma(t) > 0$ we have $|\mathcal{D}_t| \leq \alpha \cdot n \cdot \log_2(\log_2 n)/\log_2 n$.*

Proof. If $P_\sigma(t) > 0$, then there exists a constant β such that t and each of its subtrees is β-depth-balanced. By Lemma 5 this implies that every subtree t' has depth at most $c \cdot \log_2 |t'|$ for a constant c that only depends on σ. By [5, Theorem 12] (which was implicitly shown in [6]) it follows that there exists a constant α (again, only dependent on σ) such that $|\mathcal{D}_t| \leq \alpha \cdot n \cdot \log_2(\log_2 n)/\log_2 n$. □

Corollary 2. *Let $\sigma \in \Sigma_{depth}$ be a depth-balanced mapping and let $\mathcal{S} = ((\mathcal{T}^i)_{i\geq 0}, P_\sigma)$ be the corresponding leaf-centric tree source. Then, we have $R(E_{dag}, \mathcal{S}, i) \leq O((\log\log i)^2/\log i)$.*

Proof. Let α be the constant from Lemma 6. Let $t \in \mathcal{T}^i$ such that $P_\sigma(t) > 0$. Lemma 1 and 4 imply

$$\frac{1}{|t|} \cdot (|E_{\mathrm{dag}}(t)| + \log_2(P_\sigma(t))) \leq O(|\mathcal{D}_t|/|t|) + O(|\mathcal{D}_t|/|t| \cdot \log_2(|t|/|\mathcal{D}_t|)).$$

Consider the mapping $g(x) = x \cdot \log_2(1/x)$. It is monotonically increasing for $0 \leq x \leq 1/e$. Note that for all $t \in \mathcal{T}^i$ we have $|t| \geq i + 1$. Hence, if i is large enough, then Lemma 6 yields for all $t \in \mathcal{T}^i$ with $P_\sigma(t) > 0$ that

$$|\mathcal{D}_t|/|t| \leq \alpha \cdot \log_2(\log_2 |t|)/\log_2 |t| \leq \log_2(\log_2(i+1))/\log_2(i+1) \leq 1/e.$$

We obtain

$$\frac{1}{|t|} \cdot (|E_{\mathrm{dag}}(t)| + \log_2(P_\sigma(t))) \leq O((\log\log i)^2/\log i).$$

This proves the corollary. □

4 Tree compression with TSLPs

In this section, we will use general TSLPs for the compression of binary trees. The limitations of DAGs for universal source coding can be best seen for a tree source $((\mathcal{F}_i)_{i\in\mathbb{N}}, P)$ such that $P(t) > 0$ for all $t \in \mathcal{T}$. Example 2 shows that for every $n \geq 1$, there is a tree $t \in \mathcal{T}_n$ with $|\mathcal{D}_t| = n$. In that case, the bound stated in Lemma 1 cannot be used to show that the worst-case redundancy converges to zero.

4.1 TSLPs in normal form

In this section, we will use TSLPs in a certain normal form, which we introduce first.

A TSLP $\mathcal{G} = (V, A_0, r)$ is in *normal form* if the following conditions hold:

- $V = \{A_0, A_1, \ldots, A_{n-1}\}$ for some $n \in \mathbb{N}$, $n \geq 1$.
- For every $A_i \in V_0$, the right-hand side $r(A_i)$ is a term of the form $A_j(\alpha)$, where $A_j \in V_1$ and $\alpha \in V_0 \cup \{a\}$.
- For every $A_i \in V_1$ the right-hand side $r(A_i)$ is a term of the form $A_j(A_k(x))$, $f(\alpha, x)$, or $f(x, \alpha)$, where $A_j, A_k \in V_1$ and $\alpha \in V_0 \cup \{a\}$.
- For every $A_i \in V$ define the word $\rho(A_i) \in (V \cup \{a\})^*$ as follows:

$$\rho(A_i) = \begin{cases} A_j\alpha & \text{if } r(A_i) = A_j(\alpha) \\ A_jA_k & \text{if } r(A_i) = A_j(A_k(x)) \\ \alpha & \text{if } r(A_i) = f(\alpha, x) \text{ or } f(x, \alpha) \end{cases}$$

Let $\rho_{\mathcal{G}} = \rho(A_0)\rho(A_1)\cdots\rho(A_{n-1}) \in \{a, A_1, A_2, \ldots, A_{n-1}\}^*$. Then we require that $\rho_{\mathcal{G}}$ is of the form $\rho_{\mathcal{G}} = A_1u_1A_2u_2\cdots A_{n-1}u_{n-1}$ with $u_i \in \{a, A_1, A_2, \ldots, A_i\}^*$.

- $\mathsf{val}_{\mathcal{G}}(A_i) \neq \mathsf{val}_{\mathcal{G}}(A_j)$ for $i \neq j$

As for DAGs we also allow the TSLP $\mathcal{G}_a = (\{A_0\}, A_0, A_0 \mapsto a)$ in order to get the singleton tree a. In this case, we set $\rho_{\mathcal{G}_a} = \rho(A_0) = a$.

Let $\mathcal{G} = (V, A_0, r)$ be a TSLP in normal form with $V = \{A_0, A_1, \ldots, A_{n-1}\}$ for the further definitions. We define the size of $\mathcal{G}$ as $|\mathcal{G}| = |\rho_{\mathcal{G}}|$. This is the total number of occurrences of symbols from $V \cup \{a\}$ in all right-hand sides of $\mathcal{G}$. Let $\omega_{\mathcal{G}}$ be the word obtained from $\rho_{\mathcal{G}}$ by removing for every $1 \leq i \leq n-1$ the first occurrence of A_i from $\rho_{\mathcal{G}}$. Thus, if $\rho_{\mathcal{G}} = A_1u_1A_2u_2\cdots A_{n-1}u_{n-1}$ with $u_i \in \{a, A_1, A_2, \ldots, A_i\}^*$, then $\omega_{\mathcal{G}} = u_1u_2\cdots u_{n-1}$. The *entropy* $H(\mathcal{G})$ of the normal form TSLP $\mathcal{G}$ is defined as the empirical unnormalized entropy of the word $\omega_{\mathcal{G}}$:

$$H(\mathcal{G}) = H(\omega_{\mathcal{G}}).$$

Example 3. *Let $\mathcal{G} = (\{A_0, A_1, A_2, A_3, A_4\}, A_0, r)$ be the normal form TSLP with $A_0, A_2, A_3 \in V_0, A_1, A_4 \in V_1$ and*

$$r(A_0) = A_1(A_2),\ r(A_1) = f(x, A_3),\ r(A_2) = A_4(A_3),$$
$$r(A_3) = A_4(a),\ r(A_4) = f(x, a).$$

We have $\mathsf{val}(\mathcal{G}) = f(f(f(a, a), a), f(a, a))$, $\rho_{\mathcal{G}} = A_1A_2A_3A_4A_3A_4aa$ ($u_1 = u_2 = u_3 = \varepsilon$, $u_4 = A_3A_4aa$), $|\mathcal{G}| = 8$ and $\omega_{\mathcal{G}} = A_3A_4aa$.

The *derivation tree* $T_{\mathcal{G}}$ of $\mathcal{G}$ is a rooted tree, where every node is labelled with a symbol from $V \cup \{a\}$. The root is labelled with A_0. Nodes labelled with a are the leaves of $\mathcal{T}_G$. A node v that is labelled with a nonterminal A_i has $|\rho(A_i)|$ many children. If $\rho(A_i) = \alpha \in V_0 \cup \{a\}$ then the single child of v is labelled with α. If $\rho(A_i) = A_j\alpha$ with $\alpha \in V \cup \{a\}$ then the left (resp., right) child of v is labelled with A_j (resp., α). An *initial subtree* of $T_{\mathcal{G}}$ is a tree that can be obtained from $T_{\mathcal{G}}$ as follows: Take a subset U of the nodes of $T_{\mathcal{G}}$ and remove from $T_{\mathcal{G}}$ all proper descendants of nodes from U, i.e., all nodes that are located strictly below a node from U. For such an initial subtree we denote with $\mathsf{val}_{\mathcal{G}}(T') \in \mathcal{T}_V$ the tree that is derived from the start nonterminal A_0 by using the rules according to T'. Formally, it is defined by assigning to every node v of T' a tree or context $t_v \in \mathcal{T}_V \cup \mathcal{C}_V$ inductively as follows:

- If v is labelled with a then $t_v = a$.
- If v is labelled with $A_i \in V_0$ and is a leaf of T', then $t_v = A_i$.
- If v is labelled with $A_i \in V_1$ and is a leaf of T', then $t_v = A_i(x)$.

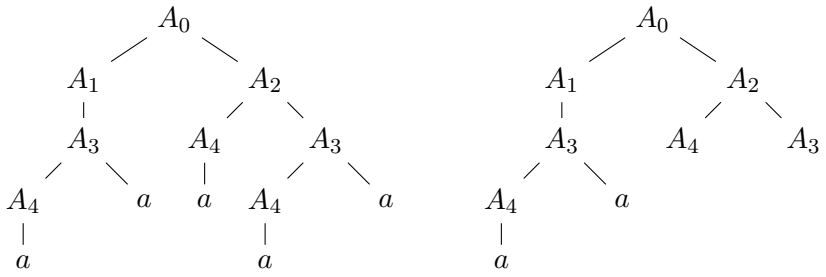

Figure 1: The derivation tree $T_{\mathcal{G}}$ of the TSLP from Example 4 (left) and an initial subtree T' of $T_{\mathcal{G}}$ (right).

- If v is labelled with $A_i \in V_1$ and has a single child node u, then $t_v = r(A_i)(t_u)$ (note that $r(A_i)$ must be of the form $f(x, \alpha)$ or $f(\alpha, x)$).
- If v has the left child u_1 and the right child u_2, then $t_v = t_{u_1}(t_{u_2})$.

We finally set $\mathsf{val}_{\mathcal{G}}(T') = t_{v_0}$, where v_0 is the root node of T'.

Example 4. *Let $\mathcal{G}$ be the normal form TSLP from Example 3. The derivation tree $T_{\mathcal{G}}$ is shown in Figure 1 on the left; an initial subtree T' of it is shown on the right. We have $\mathsf{val}_{\mathcal{G}}(T') = f(A_4(A_3), f(a, a))$.*

A *grammar-based tree compressor* is an algorithm ψ that produces for a given tree $t \in \mathcal{T}$ a TSLP $\mathcal{G}_t$ in normal form. The *compression ratio* of ψ is the mapping $n \mapsto \gamma_\psi(n)$ with

$$\gamma_\psi(n) = \max_{t \in \mathcal{T}_n} |\mathcal{G}_t|/n.$$

It is not hard to show that every TSLP can be transformed with a linear size increase into a normal form TSLP that derives the same tree. For example, the TSLP from Example 1 is transformed into the normal form TSLP described in Example 3. We will not use this fact, since all we need is the following theorem from [5]:

Theorem 1. *There exists a grammar-based compressor ψ (working in linear time) with $\gamma_\psi(n) \in O(1/\log n)$.*

4.2 Binary coding of TSLPs in normal form

In this section we fix a binary encoding for normal form TSLPs. This encoding is similar to the one for SLPs [8] and DAGs [17]. Let $\mathcal{G} = (V, A_0, r)$ be a TSLP in normal form with $n = |V|$ nonterminals. Let $m = |\mathcal{G}| = |\rho_{\mathcal{G}}|$ be the size of $\mathcal{G}$. We define the type $\mathsf{type}(A_i) \in \{0, 1, 2, 3\}$ of a nonterminal

$A_i \in V$ as follows:

$$\mathsf{type}(A_i) = \begin{cases} 0 & \text{if } \rho(A_i) \in V_1(V_0 \cup \{a\}) \\ 1 & \text{if } \rho(A_i) \in V_1V_1 \\ 2 & \text{if } \rho(A_i) = f(\alpha, x) \text{ for some } \alpha \in V_0 \cup \{a\} \\ 3 & \text{if } \rho(A_i) = f(x, \alpha) \text{ for some } \alpha \in V_0 \cup \{a\} \end{cases}$$

We define the binary word $B(\mathcal{G}) = w_0w_1w_2w_3w_4$, where the words $w_i \in \{0,1\}^+$, $0 \le i \le 4$, are defined as follows:

- $w_0 = 0^{n-1}1$,
- $w_1 = a_0b_0a_1b_1 \cdots a_{n-1}b_{n-1}$, where a_jb_j is the 2-bit binary encoding of $\mathsf{type}(A_j)$,
- Let $\rho_{\mathcal{G}} = A_1u_1A_2u_2 \cdots A_{n-1}u_{n-1}$ with $u_i \in \{a, A_1, A_2, \ldots, A_i\}^*$. Then $w_2 = 10^{|u_1|}10^{|u_2|} \cdots 10^{|u_{n-1}|}$.
- For $1 \le i \le n-1$ let $k_i = |\rho_{\mathcal{G}}|_{A_i} \ge 1$ be the number of occurrences of the nonterminal A_i in the word $\rho_{\mathcal{G}}$. Then $w_3 = 0^{k_1-1}10^{k_2-1}1 \cdots 0^{k_{n-1}-1}1$.
- The word w_4 encodes the word $\omega_{\mathcal{G}}$ using the well-known enumerative encoding [3]. Every nonterminal A_i, $1 \le i \le n-1$, has $\eta(A_i) := k_i - 1$ occurrences in $\omega_{\mathcal{G}}$. The symbol a has $\eta(a) := m - (k_1 + \cdots + k_{n-1})$ many occurrences in $\omega_{\mathcal{G}}$. Let S be the set of words over the alphabet $\{a, A_1, \ldots, A_{n-1}\}$ with $\eta(a)$ occurrences of a and $\eta(A_i)$ occurrences of A_i for every $1 \le i \le n-1$. Hence,

$$|S| = \frac{(m-n+1)!}{\eta(a)! \cdot \prod_{i=1}^{n-1} \eta(A_i)!}. \tag{9}$$

  Let $v_0, v_1, \ldots, v_{|S|-1}$ be the lexicographic enumeration of the words from S with respect to the alphabet order $a, A_1, \ldots, A_{n-1}$. Then w_4 is the binary encoding of the unique index i such that $\omega_{\mathcal{G}} = v_i$, where $|w_4| = \lceil \log_2 |S| \rceil$ (leading zeros are added to the binary encoding of i to obtain the length $\lceil \log_2 |S| \rceil$).

Example 5. *Consider the normal from TSLP $\mathcal{G}$ from Example 3. We have $w_0 = 00001$, $w_1 = 0011000011$, $w_2 = 11110000$ and $w_3 = 110101$. To compute w_4, note first that there are $|S| = 12$ words with two occurrences of a and one occurrence of A_3 and A_4. It follows that $|w_4| = \lceil \log_2(12) \rceil = 4$. Further, since the order of the alphabet is a, A_3, A_4, there are only three words in S (A_4A_3aa, A_4aA_3a and A_4aaA_3), which are lexicographically larger than $\omega_{\mathcal{G}} = A_3A_4aa$. Hence, $\omega_{\mathcal{G}} = v_8$ and thus $w_4 = 1000$.*

Lemma 7. *The set of code words $B(\mathcal{G})$, where $\mathcal{G}$ ranges over all TSLPs in normal form, is a prefix code.*

Proof. Let $B(\mathcal{G}) = w_0w_1w_2w_3w_4$ with w_i defined as above. We show how to recover the TSLP $\mathcal{G}$. From w_0 we can determine $n = |V|$ and the factor w_1 of $B(\mathcal{G})$. Hence, we can determine the type of every nonterminal. Using the types, we can compute $|\mathcal{G}| = |\rho_{\mathcal{G}}|$. Moreover, the types allow to compute $\mathcal{G}$ from the word ρ_G. Hence, it remains to determine $\rho_{\mathcal{G}}$. From $|\mathcal{G}|$ one can compute w_2. To compute $\rho_{\mathcal{G}}$ from w_2, one only needs $\omega_{\mathcal{G}}$. For this, one determines the frequencies $\eta(a), \eta(A_1), \ldots, \eta(A_{n-1})$ of the symbols in $\omega_{\mathcal{G}}$ from w_3. Using these frequencies one computes the size $|S|$ from (9) and the length $\lceil \log_2 |S| \rceil$ of w_4. From w_4, one can finally compute ω_G. □

Note that $|B(\mathcal{G})| \leq O(|\mathcal{G}|) + |w_4|$. By using the well-known bound on the code length of enumerative encoding [4, Theorem 11.1.3], we get:

Lemma 8. *For the length of the binary coding $B(\mathcal{G})$ we have: $|B(\mathcal{G})| \leq O(|\mathcal{G}|) + H(\mathcal{G})$.*

4.3 Universal source coding based on TSLPs in normal form

Let $((\mathcal{F}_i)_{i\in\mathbb{N}}, P)$ be a tree source as defined in Section 2.2. We say that $((\mathcal{F}_i)_{i\in\mathbb{N}}, P)$ has the *strong domination property* if the there exists a mapping $\lambda : \mathcal{T} \cup \mathcal{C} \to \mathbb{R}_{>0}$ with the following properties:

(i) $\lambda(t) \geq P(t)$ for every $t \in \mathcal{T}$.

(ii) $\lambda(f(s,t)) \leq \lambda(s) \cdot \lambda(t)$ for all $s, t \in \mathcal{T}$

(iii) $\lambda(s(t)) \leq \lambda(s) \cdot \lambda(t)$ for all $s \in \mathcal{C}$ and $t \in \mathcal{T}$

(iv) There are constants c_1, c_2 such that $\sum_{t\in\mathcal{T}_n\cup\mathcal{C}_n} \lambda(t) \leq c_1 \cdot n^{c_2}$ for all $n \geq 1$.

The proof of the following lemma combines ideas from [8] and [17].

Lemma 9. *Assume that $((\mathcal{F}_i)_{i\in\mathbb{N}}, P)$ has the strong domination property. Let $t \in \mathcal{T}_n$ with $n \geq 2$ and $P(t) > 0$, and let $\mathcal{G} = (V, A_0, r)$ be a TSLP in normal form with $\mathsf{val}(\mathcal{G}) = t$. We have*

$$H(G) \leq -\log_2 P(t) + O(|\mathcal{G}|) + O(|\mathcal{G}| \cdot \log_2(n/|\mathcal{G}|)).$$

Proof. Let $m = |\mathcal{G}| = |\rho_{\mathcal{G}}|$ be the size of $\mathcal{G}$, $k = |V|$, and $\ell := m+1-k \leq m$. Let $T = T_{\mathcal{G}}$ be the derivation tree of $\mathcal{G}$. We define an initial subtree T' as follows: If v_1 and v_2 are non-leaf nodes of T that are labelled with the same nonterminal and v_1 comes before v_2 in preorder, then we remove from T all proper descendants of v_2. Thus, for every $A_i \in V$ there is exactly one non-leaf node in T' that is labelled with A_i. Let $t' = \mathsf{val}_{\mathcal{G}}(T')$. In the derivation of t' from A_0, every rule of $\mathcal{G}$ is used exactly once. For the TSLP from Example 3, the tree T' is shown in Figure 1 on the right, and the tree t' is computed in Example 4.

Note that T' has exactly $m+1$ many nodes and k non-leaf nodes. Thus, T' has ℓ leaves. Let $v_1, v_2, \ldots, v_\ell$ be the sequence of all leaves of T' (w.l.o.g. in preorder) and let $\alpha_i \in \{a, A_1, \ldots, A_{k-1}\}$ be the label of v_i. Let $\overline{\alpha} = (\alpha_1, \alpha_2, \ldots, \alpha_\ell)$. Then $|\omega_{\mathcal{G}}|_\alpha = |\overline{\alpha}|_\alpha$ for every $\alpha \in \{a, A_1, \ldots, A_{k-1}\}$, and this is also the number of occurrences of α in the tree t'. Hence, $p_{\overline{\alpha}}$ and $p_{\omega_{\mathcal{G}}}$ have the same empirical distribution. Let $t_i = \mathsf{val}_{\mathcal{G}}(\alpha_i) \in \mathcal{T} \cup \mathcal{C}$. For the TSLP from Example 3 we get $\overline{\alpha} = (a, a, A_4, A_3)$. Since $\mathsf{val}_{\mathcal{G}}(A_i) \neq \mathsf{val}_{\mathcal{G}}(A_j)$ for all $i \neq j$ and $\mathsf{val}_{\mathcal{G}}(A_i) \neq a$ for all i (this holds for every normal form TSLP that produces a tree of size at least two), the tuple $\overline{t} = (t_1, t_2, \ldots, t_\ell)$ satisfies $p_{\omega_{\mathcal{G}}}(\alpha_i) = p_{\overline{t}}(t_i)$ for all $1 \leq i \leq \ell$.

Using conditions (ii) and (iii) of the strong domination property, we get

$$\lambda(t) \leq \prod_{i=1}^{\ell} \lambda(t_i). \tag{10}$$

For $j \in \mathbb{N}$, $j \geq 1$, let

$$M_j = \sum_{u \in \mathcal{T}_j \cup \mathcal{C}_j} \lambda(u) \leq c_1 \cdot j^{c_2},$$

where c_1 and c_2 are the constants from condition (iv) of the strong domination property. Let $D := 6/\pi^2 \geq 1/2$ and define for every $u \in \mathcal{T}_j \cup \mathcal{C}_j$:

$$q(u) := \frac{D \cdot \lambda(u)}{M_j \cdot j^2} > 0. \tag{11}$$

We get

$$\sum_{j \geq 1} \sum_{u \in \mathcal{T}_j \cup \mathcal{C}_j} q(u) = D \cdot \sum_{j \geq 1} \frac{1}{j^2} = 1.$$

Hence, we have $q(t_1) + q(t_2) + \cdots + q(t_\ell) \leq 1$. Using Shannon's inequality (2) we get

$$H(G) = H(\omega_{\mathcal{G}}) = \sum_{i=1}^{\ell} -\log_2 p_{\omega_{\mathcal{G}}}(\alpha_i) = \sum_{i=1}^{\ell} -\log_2 p_{\overline{t}}(t_i) \leq \sum_{i=1}^{\ell} -\log_2 q(t_i).$$

Using (11) and $D \geq 1/2$ we obtain

$$\begin{aligned}
H(G) &\leq \sum_{i=1}^{\ell} -\log_2 \left(\frac{D \cdot \lambda(t_i)}{M_{|t_i|} \cdot |t_i|^2} \right) \\
&= -\ell \cdot \log_2 D - \sum_{i=1}^{\ell} \log_2 \lambda(t_i) + \sum_{i=1}^{\ell} \log_2 M_{|t_i|} + 2 \sum_{i=1}^{\ell} \log_2 |t_i| \\
&\leq \ell - \sum_{i=1}^{\ell} \log_2 \lambda(t_i) + \sum_{i=1}^{\ell} (\log_2 c_1 + c_2 \log_2 |t_i|) + 2 \sum_{i=1}^{\ell} \log_2 |t_i| \\
&= (1 + \log_2 c_1) \cdot \ell - \sum_{i=1}^{\ell} \log_2 \lambda(t_i) + (2 + c_2) \cdot \sum_{i=1}^{\ell} \log_2 |t_i|.
\end{aligned}$$

From (10) and condition (i) of the strong domination property we get

$$\sum_{i=1}^{\ell} \log_2 \lambda(t_i) \geq \log_2 \lambda(t) \geq \log_2 P(t).$$

Moreover, Jensen's inequality gives

$$\sum_{i=1}^{\ell} \log_2 |t_i| \leq \ell \cdot \log_2 \left(\frac{1}{\ell} \cdot \sum_{i=1}^{\ell} |t_i| \right) = \ell \cdot \log_2(n/\ell).$$

With $\ell \leq m$ we obtain

$$\begin{aligned} H(G) &\leq (1 + \log_2 c_1) \cdot \ell - \log_2 P(t) + (2 + c_2) \cdot \ell \cdot \log_2(n/\ell) \\ &\leq -\log_2 P(t) + (1 + \log_2 c_1) \cdot m + (2 + c_2) \cdot m \cdot \log_2(n/m). \end{aligned}$$

This shows the lemma. □

Let $\psi : t \mapsto \mathcal{G}_t$ be a grammar-based tree compressor. We then consider the tree encoder $E_\psi : \mathcal{T} \to \{0,1\}^*$ defined by $E_\psi(t) = B(\mathcal{G}_t)$. Recall the definition of the worst-case redundancy $R(E_\psi, \mathcal{S}, i)$ from Section 2.2.

Theorem 2. *Assume that $\mathcal{S} = ((\mathcal{F}_i)_{i\in\mathbb{N}}, P)$ has the strong domination property. Let ψ be a grammar-based compressor such that $\gamma_\psi(n) \leq \gamma(n)$ for a monotonically decreasing function $\gamma(n)$ with $\lim_{n\to\infty} \gamma(n) = 0$. Let $n_i = \min\{|t| \mid t \in \mathcal{F}_i\}$ and assume that $n_i < n_{i+1}$ for all $i \in \mathbb{N}$.[1] Then, we have*

$$R(E_\psi, \mathcal{S}, i) \leq O(\gamma(n_i) \cdot \log_2(1/\gamma(n_i))).$$

Proof. Let $t \in \mathcal{T}$ and $\psi(t) = \mathcal{G}_t$. With Lemma 8 and 9 we get

$$\begin{aligned} \frac{1}{|t|} \cdot (|B(\mathcal{G}_t)| + \log_2 P(t)) &\leq \frac{1}{|t|} \cdot (H(\mathcal{G}_t) + O(|\mathcal{G}_t|) + \log_2 P(t)) \\ &\leq \frac{1}{|t|} \cdot (O(|\mathcal{G}_t|) + O(|\mathcal{G}_t| \cdot \log_2(|t|/|\mathcal{G}_t|)) \\ &= O(|\mathcal{G}_t|/|t|) + O(|\mathcal{G}_t|/|t| \cdot \log_2(|t|/|\mathcal{G}_t|)). \end{aligned}$$

Consider the mapping g with $g(x) = x \cdot \log_2(1/x)$. It is monotonically increasing for $0 \leq x \leq 1/e$. If i is large enough, we have for all $t \in \mathcal{F}_i$ that

$$|\mathcal{G}_t|/|t| \leq \gamma_\psi(|t|) \leq \gamma(|t|) \leq \gamma(n_i) \leq 1/e.$$

Hence, we get

$$|\mathcal{G}_t|/|t| \cdot \log_2(|t|/|\mathcal{G}_t|) = g(|\mathcal{G}_t|/|t|) \leq g(\gamma(n_i)) = \gamma(n_i) \cdot \log_2(1/\gamma(n_i)).$$

[1] This is the case for leaf-centric and depth-centric tree sources

This implies

$$\begin{aligned} R(E_\psi, \mathcal{S}, i) &= \max_{t \in \mathcal{F}_i, P(t)>0} \frac{1}{|t|} \cdot (|B(\mathcal{G}_t)| + \log_2 P(t)) \\ &\leq O(\gamma(n_i)) + O(\gamma(n_i) \cdot \log(1/\gamma(n_i))) \\ &= O(\gamma(n_i) \cdot \log(1/\gamma(n_i))), \end{aligned}$$

which proves the theorem. □

Note that the minimal size of a tree in $\mathcal{T}_{i+1}$ (resp. $\mathcal{T}^i$) is $i+1$. Hence, Theorem 1 and 2 yield:

Corollary 3. *There exists a grammar-based tree compressor ψ (working in linear time) such that $R(E_\psi, \mathcal{S}, i) \leq O(\log\log i / \log i)$ for every leaf-centric or depth-centric tree source $\mathcal{S}$ having the strong domination property.*

In the rest of the paper, we will present classes of leaf-centric and depth-centric tree sources that have the strong domination property.

4.4 Leaf-centric binary tree sources

Recall the definition of the class of mappings Σ_{leaf} by equations (3), (4), and (5) in Section 3.3 and the corresponding class of leaf-centric tree sources. In this section, we state a condition on the mapping $\sigma \in \Sigma_{\text{leaf}}$ that enforces the strong domination property for the leaf-centric tree source $((\mathcal{T}_i)_{i\geq 1}, P_\sigma)$. This allows to apply Corollary 3.

Theorem 3. *If $\sigma \in \Sigma_{leaf}$ satisfies $\sigma(i,j) \geq \sigma(i,j+1)$ and $\sigma(i,j) \geq \sigma(i+1,j)$ for all $i,j \geq 1$, then $((\mathcal{T}_i)_{i\geq 1}, P_\sigma)$ has the strong domination property.*

Proof. First, we naturally extend P_σ to a context $t \in \mathcal{C}$ using equations (4) and (5), where we set $\sigma(0,k) = \sigma(k,0) = 1$ for all $k \geq 1$ and $P_\sigma(x) = 1$. Note that $|x| = 0$. Also note that P_σ is not a probability distribution on $\mathcal{C}_n$. We have for instance $\sum_{t\in\mathcal{C}_1} P_\sigma(t) = P_\sigma(f(x,a)) + P_\sigma(f(a,x)) = 2$. We denote by C_n the n^{th} Catalan number. It is well-known that $|\mathcal{T}_n| = C_n$ for all $n \geq 1$. We set $C_0 = 1$ and define $\lambda : \mathcal{T} \cup \mathcal{C} \to \mathbb{R}_{>0}$ by

$$\lambda(t) = \max\left\{\frac{1}{C_{|t|}}, P_\sigma(t)\right\}.$$

We show the four points from the strong domination property for the mapping λ.

The first point of the strong domination property, i.e., $\lambda(t) \geq P_\sigma(t)$ for all $t \in \mathcal{T}$, is obviously true. We now prove the second point, i.e., $\lambda(f(s,t)) \leq \lambda(s) \cdot \lambda(t)$ for all $s,t \in \mathcal{T}$. Assume first that $\lambda(f(s,t)) = 1/C_{|s|+|t|}$. The inequality $C_{m+k} \geq C_m \cdot C_k$ for all $m,k \geq 0$ yields

$$\frac{1}{C_{|s|+|t|}} \leq \frac{1}{C_{|s|}} \cdot \frac{1}{C_{|t|}} \leq \lambda(s) \cdot \lambda(t). \tag{12}$$

On the other hand, if $\lambda(f(s,t)) = P_\sigma(f(s,t))$, then we have

$$P_\sigma(f(s,t)) = \sigma(|s|,|t|) \cdot P_\sigma(s) \cdot P_\sigma(t) \leq P_\sigma(s) \cdot P_\sigma(t) \leq \lambda(s) \cdot \lambda(t),$$

since $0 \leq \sigma(i,j) \leq 1$ for all i,j.

We now consider the third point, i.e., $\lambda(s(t)) \leq \lambda(s) \cdot \lambda(t)$ for all $s \in \mathcal{C}$ and $t \in \mathcal{T}$. Note first that the case $\lambda(s(t)) = 1/C_{|s|+|t|}$ follows again from equation (12), since $|s(t)| = |s| + |t|$. So we assume that $\lambda(s(t)) = P_\sigma(s(t))$. Let k be the length (measured in the number of edges) from the root of s to the unique x-labelled node in s.

We show $P_\sigma(s(t)) \leq P_\sigma(s) \cdot P_\sigma(t)$ by induction over $k \geq 0$. If $k = 0$ then $s = x$ and we get $P_\sigma(s(t)) = P_\sigma(t) = P_\sigma(x) \cdot P_\sigma(t) = P_\sigma(s) \cdot P_\sigma(t)$. Let us now assume that $k \geq 1$. Then, s must have the form $s = f(u,v)$. Without loss of generality assume that x occurs in u; the other case is of course symmetric. Therefore, we have $s(t) = f(u(t),v)$. The tree $u(t)$ fulfills the induction hypothesis and therefore $P_\sigma(u(t)) \leq P_\sigma(u) \cdot P_\sigma(t)$. Moreover, we have $\sigma(|u(t)|,|v|) = \sigma(|u|+|t|,|v|) \leq \sigma(|u|,|v|)$. We get

$$\begin{aligned} P_\sigma(s(t)) = P_\sigma(f(u(t),v)) &= P_\sigma(u(t)) \cdot \sigma(|u(t)|,|v|) \cdot P_\sigma(v) \\ &\leq P_\sigma(t) \cdot P_\sigma(u) \cdot \sigma(|u|,|v|) \cdot P_\sigma(v) \\ &= P_\sigma(t) \cdot P_\sigma(s). \end{aligned}$$

For the fourth property we show $\sum_{t \in \mathcal{T}_n \cup \mathcal{C}_n} \lambda(t) \leq 8n-2$ for all $n \geq 1$. First, we have

$$\sum_{t \in \mathcal{T}_n} \lambda(t) \leq \sum_{t \in \mathcal{T}_n} \left(C_n^{-1} + P_\sigma(t)\right) = 2. \tag{13}$$

We show that $\sum_{t \in \mathcal{C}_n} \lambda(t) \leq 8n-4$. Every tree $t \in \mathcal{T}_n$ has $2n-1$ nodes. Let v be a node of t. We obtain two contexts from t and v by replacing in t the subtree $t[v]$ by either $f(x,t[v])$ or $f(t[v],x)$. Let us denote the resulting contexts by $t_{v,1}$ and $t_{v,2}$. We have $\lambda(t_{v,1}) = \lambda(t_{v,2}) = \lambda(t)$ for every node v of t. Moreover, for every context $t \in \mathcal{C}_n$ there exists a tree $t' \in \mathcal{T}_n$ and a node v of t' such that $t'_{v,1} = t$ or $t'_{v,2} = t$ (depending on whether x is the left or right child of its parent node in t). Hence, we get

$$\sum_{t \in \mathcal{C}_n} \lambda(t) \leq (4n-2) \cdot \sum_{t \in \mathcal{T}_n} \lambda(t). \tag{14}$$

Together with equation (13) we have $\sum_{t \in \mathcal{T}_n \cup \mathcal{C}_n} \lambda(t) \leq 8n-2$. □

Example 6. *An example for a leaf-centric tree source $((\mathcal{T}_i)_{i \geq 1}, P_\sigma)$, where $\sigma \in \Sigma_{leaf}$ satisfies $\sigma(i,j) \geq \sigma(i,j+1)$ and $\sigma(i,j) \geq \sigma(i+1,j)$ for all $i,j \geq 1$, is the famous* binary search tree model*; see [11] for an investigation in the context of information theory. It is obtained by setting $\sigma(i,j) = 1/(i+j)$.*

4.5 Depth-centric binary tree sources

Recall the definition of the class of mappings Σ_{depth} by equations (6), (7), and (8) in Section 3.4, and the corresponding class of depth-centric tree sources. In this section, we state a condition on the mapping $\sigma \in \Sigma_{\text{depth}}$ that enforces the strong domination property for the depth-centric tree source $((\mathcal{T}_i)_{i\geq 1}, P_\sigma)$. This allows again to apply Corollary 3.

Theorem 4. *If $\sigma \in \Sigma_{depth}$ satisfies $\sigma(i,j) \geq \sigma(i,j+1)$ and $\sigma(i,j) \geq \sigma(i+1,j)$ for all $i,j \geq 0$, then $((\mathcal{T}^i)_{i\geq 0}, P_\sigma)$ has the strong domination property.*

Proof. Recall that the depth of a context $t \in \mathcal{C}$ is defined as the depth of the tree $t(a)$. Using this information, we extend P_σ to a context $t \in \mathcal{C}$ using equations (7) and (8), where we set $P_\sigma(x) = 1$. Similarly to the proof of Theorem 3 for leaf-centric tree sources, we define

$$\lambda(t) = \max\left\{\frac{1}{C_{|t|}}, P_\sigma(t)\right\}.$$

The first point of the strong domination property, i.e., $\lambda(t) \geq P_\sigma(t)$ for all $t \in \mathcal{T}$, follows directly from the definition of λ. Now we prove the second point, i.e., $\lambda(f(s,t)) \leq \lambda(s)\cdot\lambda(t)$ for all $s,t \in \mathcal{T}$. The case $\lambda(f(s,t)) = 1/C_{|s|+|t|}$ is covered by equation (12). If otherwise $\lambda(f(s,t)) = P_\sigma(f(s,t))$, then

$$P_\sigma(f(s,t)) = \sigma(d(s),d(t))\cdot P_\sigma(s)\cdot P_\sigma(t) \leq P_\sigma(s)\cdot P_\sigma(t) \leq \lambda(s)\cdot\lambda(t),$$

since $0 \leq \sigma(i,j) \leq 1$ for all i,j.

Consider now the third point, i.e., $\lambda(s(t)) \leq \lambda(s)\cdot\lambda(t)$ for all $s \in \mathcal{C}$ and $t \in \mathcal{T}$. Note that the case $\lambda(s(t)) = 1/C_{|s|+|t|}$ follows again from equation (12). Hence, we can assume that $\lambda(s(t)) = P_\sigma(s(t))$. Again, we prove $P_\sigma(s(t)) \leq P_\sigma(s)\cdot P_\sigma(t)$ by induction over the length $k \geq 0$ (measured in the number of edges) from the root of s to the unique x-labelled node in s. If $k = 0$ then $s = x$ and $P_\sigma(s) = 1$, which gives us $P_\sigma(s(t)) = P_\sigma(t) = P_\sigma(s)\cdot P_\sigma(t)$. We now assume $k \geq 1$ and $s = f(u,v)$. Without loss of generality assume that x occurs in u; the other case is symmetric. Therefore, we have $s(t) = f(u(t),v)$. We apply the induction hypothesis to the tree $u(t)$, which yields $P_\sigma(u(t)) \leq P_\sigma(u)\cdot P_\sigma(t)$. Moreover, since $d(u) \leq d(u(t))$ we have $\sigma(d(u(t)),d(v)) \leq \sigma(d(u),d(v))$. It follows that

$$\begin{aligned} P_\sigma(s(t)) = P_\sigma(f(u(t),v)) &= P_\sigma(u(t))\cdot\sigma(d(u(t)),d(v))\cdot P_\sigma(v) \\ &\leq P_\sigma(t)\cdot P_\sigma(u)\cdot\sigma(d(u),d(v))\cdot P_\sigma(v) \\ &= P_\sigma(t)\cdot P_\sigma(s). \end{aligned}$$

For the fourth property we show $\sum_{t\in\mathcal{T}_n\cup\mathcal{C}_n}\lambda(t) \leq 4n^2+3n-1$ for all $n \geq 1$. First, we have

$$\sum_{t\in\mathcal{T}_n}\lambda(t) \leq \sum_{t\in\mathcal{T}_n}\left(C_n^{-1} + P_\sigma(t)\right) = 1 + \sum_{t\in\mathcal{T}_n} P_\sigma(t). \tag{15}$$

Note that P_σ is not a probability distribution on $\mathcal{T}_n$ since this section deals with depth-centric tree sources. But for each tree $t \in \mathcal{T}_n$ we have $\lceil \log_2(n) \rceil \leq d(t) \leq n-1$, which yields

$$\sum_{t \in \mathcal{T}_n} P_\sigma(t) \leq \sum_{i=\lceil \log_2(n) \rceil}^{n-1} \sum_{t \in \mathcal{T}^i} P_\sigma(t) = n - \lceil \log_2(n) \rceil \leq n.$$

Together with equation (15) we get $\sum_{t \in \mathcal{T}_n} \lambda(t) \leq n+1$. The remaining part $\sum_{t \in \mathcal{C}_n} \lambda(t)$ can be estimated with help of equation (14) from the corresponding part in the proof of Theorem 3. In total, we have

$$\sum_{t \in \mathcal{T}_n \cup \mathcal{C}_n} \lambda(t) \leq (4n-1) \sum_{t \in \mathcal{T}_n} \lambda(t) \leq (4n-1)(n+1) = 4n^2 + 3n - 1.$$

□

5 Future research

We plan to investigate, whether the strong domination property can be shown also for other classes of tree sources. An interesting class are the tree sources derived from stochastic context-free grammars [15]. Another interesting question is, whether convergence rate of $O(\log \log i / \log i)$ in Corollary 3 can be improved to $O(1/\log i)$. In the context of grammar-based string compression, such an improvement has been accomplished in [9].

References

[1] Janos Aczél. On Shannon's inequality, optimal coding, and characterizations of Shannon's and Renyi's entropies. Technical Report Research Report AA-73-05, University of Waterloo, 1973. `https://cs.uwaterloo.ca/research/tr/1973/CS-73-05.pdf`.

[2] Yongwook Choi and Wojciech Szpankowski. Compression of graphical structures: Fundamental limits, algorithms, and experiments. *IEEE Transactions on Information Theory*, 58(2):620–638, 2012.

[3] Thomas M. Cover. Enumerative source encoding. *IEEE Transactions on Information Theory*, 19(1):73–77, 1973.

[4] Thomas M. Cover and Joy A. Thomas. *Elements of information theory (2. ed.)*. Wiley, 2006.

[5] Moses Ganardi, Danny Hucke, Artur Jeż, Markus Lohrey, and Eric Noeth. Constructing small tree grammars and small circuits for formulas. *Journal of Computer and System Sciences*, 2016. `http://dx.doi.org/10.1016/j.jcss.2016.12.007`.

[6] Lorenz Hübschle-Schneider and Rajeev Raman. Tree compression with top trees revisited. In *Proceedings of the 14th International Symposium on Experimental Algorithms, SEA 2015*, volume 9125 of *Lecture Notes in Computer Science*, pages 15–27. Springer, 2015.

[7] John C. Kieffer. A survey of Bratteli information source theory. In *Proceedings of the IEEE International Symposium on Information Theory, ISIT 2016*, pages 16–20. IEEE, 2016.

[8] John C. Kieffer and En hui Yang. Grammar-based codes: A new class of universal lossless source codes. *IEEE Transactions on Information Theory*, 46(3):737–754, 2000.

[9] John C. Kieffer and En hui Yang. Structured grammar-based codes for universal lossless data compression. *Communications in Information and Systems*, 2(1):29–52, 2002.

[10] John C. Kieffer, En-Hui Yang, Gregory J. Nelson, and Pamela C. Cosman. Universal lossless compression via multilevel pattern matching. *IEEE Transactions on Information Theory*, 46(4):1227–1245, 2000.

[11] John C. Kieffer, En-Hui Yang, and Wojciech Szpankowski. Structural complexity of random binary trees. In *Proceedings of the IEEE International Symposium on Information Theory, ISIT 2009*, pages 635–639. IEEE, 2009.

[12] Markus Lohrey. Algorithmics on SLP-compressed strings: A survey. *Groups Complexity Cryptology*, 4(2):241–299, 2012.

[13] Markus Lohrey. Grammar-based tree compression. In *Proceedings of the 19th International Conference on Developments in Language Theory, DLT 2015*, volume 9168 of *Lecture Notes in Computer Science*, pages 46–57. Springer, 2015.

[14] Abram Magner, Krzysztof Turowski, and Wojciech Szpankowski. Lossless compression of binary trees with correlated vertex names. In *Proceedings of the IEEE International Symposium on Information Theory, ISIT 2016*, pages 1217–1221. IEEE, 2016.

[15] Michael I. Miller and Joseph A. O'Sullivan. Entropies and combinatorics of random branching processes and context-free languages. *IEEE Transactions Information Theory*, 38(4):1292–1310, 1992.

[16] En-Hui Yang and John C. Kieffer. Efficient universal lossless data compression algorithms based on a greedy sequential grammar transform - part one: Without context models. *IEEE Transactions on Information Theory*, 46(3):755–777, 2000.

[17] Jie Zhang, En-Hui Yang, and John C. Kieffer. A universal grammar-based code for lossless compression of binary trees. *IEEE Transactions on Information Theory*, 60(3):1373–1386, 2014.

[18] Jacob Ziv and Abraham Lempel. Compression of individual sequences via variable-rate coding. *IEEE Transactions on Information Theory*, 24(5):530–536, 1977.